T.V.'s "Jeopardy!" : A Rich Empirical Data Set for Behavioral Economics

Matthew P. Haney (JHU)

Extended Abstract

In the final round of the game show "Jeopardy!" the three contestants must make a strategic wager under conditions of uncertainty. Analyzing how real people handle this surprisingly complex strategic game, provides a novel source of data for the study of how actual human behavior deviates from the rational strategies dictated by game theory. The 900 games in this paper's data set reveal scenarios in which one or more of the players wager in a consistently sub-optimal manner, resulting in steady equilibriums that simply should not exist under the assumptions of traditional game theory. This paper indicates that the seemingly irrational or inconsistent behavior exhibited by some of the contestants is explainable using ideas from behavioral economics, specifically heuristic (rule-based) reasoning and overconfidence.

The Game

The rules of the game are extremely simple. Each of the three contestants enters Final Jeopardy (FJ) with a score (dollars) total based on the previous two rounds. If a player does not have a positive score, then he or she does not play in FJ. Each contestant knows his own score and that of his opponents. The players are shown the category of the FJ question and are then asked to make a wager between zero and their current score. If the FJ question is answered correctly, the wager is added to their score, if not it is deducted. The player with the highest score at the end of FJ gets to keep his winnings, with a minimum of $5000, and is invited back to play again, with no limit on how long a winning streak can last. The player in second place gets $2000, and third gets $1000.

The Wager Space

The entire Final Jeopardy wagering problem can be viewed as a three player non-zero sum matrix game, where each wager amount corresponds to a strategy, and the payoff values of the matrix are equal to the value of winning multiplied by the probability of winning. Player one is the player with the largest score (P_1) entering FJ, while player two has the second largest score (P_2), and player three has the smallest score (P3).

While the three player wager space is a 3-dimensional box, the simplified two dimensional version is simply a rectangle, which can be thought of as a nearly continuous payoff matrix. Below is an example of such a matrix for a game where P_1=10,000 and P2=7,000.
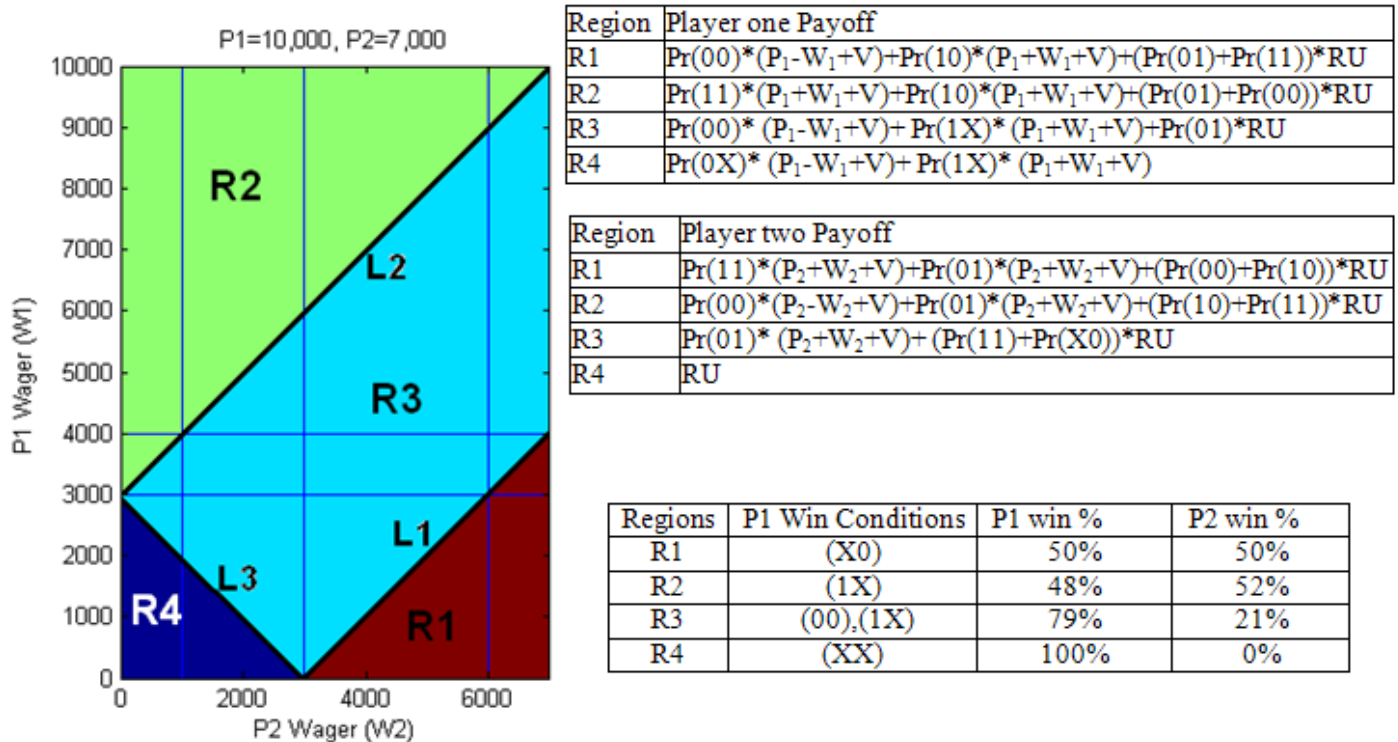

Region	Player one Payoff
R1	$Pr(00)*(P_1-W_1+V)+Pr(10)*(P_1+W_1+V)+(Pr(01)+Pr(11))*RU$
R2	$Pr(11)*(P_1+W_1+V)+Pr(10)*(P_1+W_1+V)+(Pr(01)+Pr(00))*RU$
R3	$Pr(00)*(P_1-W_1+V)+Pr(1X)*(P_1+W_1+V)+Pr(01)*RU$
R4	$Pr(0X)*(P_1-W_1+V)+Pr(1X)*(P_1+W_1+V)$

Region	Player two Payoff
R1	$Pr(11)*(P_2+W_2+V)+Pr(01)*(P_2+W_2+V)+(Pr(00)+Pr(10))*RU$
R2	$Pr(00)*(P_2-W_2+V)+Pr(01)*(P_2+W_2+V)+(Pr(10)+Pr(11))*RU$
R3	$Pr(01)*(P_2+W_2+V)+(Pr(11)+Pr(X0))*RU$
R4	RU

Regions	P1 Win Conditions	P1 win %	P2 win %
R1	(X0)	50%	50%
R2	(1X)	48%	52%
R3	(00),(1X)	79%	21%
R4	(XX)	100%	0%

Figure 1: Wager Space for P_1=10,000and P_2=7,000

In the notation of the above figures and tables, W_i is the wager of the player i, 0 denotes a player answering incorrectly, 1 denotes correctly, and X means either. (01) means that player one answered incorrectly, and player two

answered correctly. $Pr(X0)$ would then be the probability that player two answered incorrectly regardless of how player one answered. RU is simply the value of being the runner-up, $2000. V is the value associated with winning other than the immediate cash prize, including the expected winnings of future games (since the winner returns) and intangibles such as bragging rights. This value is highly subjective for each individual player and plays a key role later in the full paper.

In the figure the three black lines separating the four regions are defined by the equations $L1:W1=W2-(P_1-P_2)$, $L2:W1=(P_1-P_2)+W2$, and $L3:W1=(P_1-P_2)-W1$. As the difference between P_1 and P_2 decreases, R3 and R4 shrink while the wager-space approaches a square. This simple figure reveals some of the key dynamics of the two player wagering game, especially why player two wagering nearly everything can be such a poor decision. If player one wagers above the upper blue horizontal line ($W_1>2*P_2-P_1$) when player two makes a large wager ($W_2>3*P_2-2*P_1$), then the game ends in R3, where player two only has a 21% chance of victory. Those blue lines represent some of the heuristics that prove so important to the decision making process of the players.

The Heuristics

As demonstrated by the figure and tables in the last section, the two player wagering problem is fairly complex. When formulating a wager, a player must consider how much he values winning the game and the probability of each of the four outcomes The player must then try to figure out what his opponent will do (which depends on the opponents estimation of those other quantities). The contestants on "Jeopardy!" are in a stressful environment and not allowed to use even the simplest of calculators to aid in their wager decision. By necessity they need some way to simplify their decision, and it appears that players rely on a set heuristics, or rules of thumb, to guide their wagers. The heuristic wagers are simple, such as: almost nothing (H1), small enough not to fall behind a player ("conservative wager" H2), enough to get ahead of the lead player (H3), nearly everything (H4), and enough to keep ahead of a trailing player if he wagers everything and gets it correct ("covering wager" H5). These five heuristics are readily seen in figure 1; they are the borders of the play space and the points where the boundary lines intersect the borders.

Each of these heuristics has a wager that the opponent can use to counter it. For example, the left most blue vertical line in figure 1 counters player one wagering to cover player two doubling up by playing the intersection in R2, a region much more favorable to player two than R3. And then each counter wager has its own wager, and so on, all of which would be viable strategies if one were solving the wagering game abstractly. But the data only shows a few counter strategies in use, including player two countering player one's conservative wager and covering wager (S1 and S2).

The Data

Below are normalized scatter-plots of the 900 games in the data set. The x axis is the ratio of P_2 to P_1, and the y axis is the ratio of a players wager to his current score. The relevant heuristics and strategies are plotted as well.
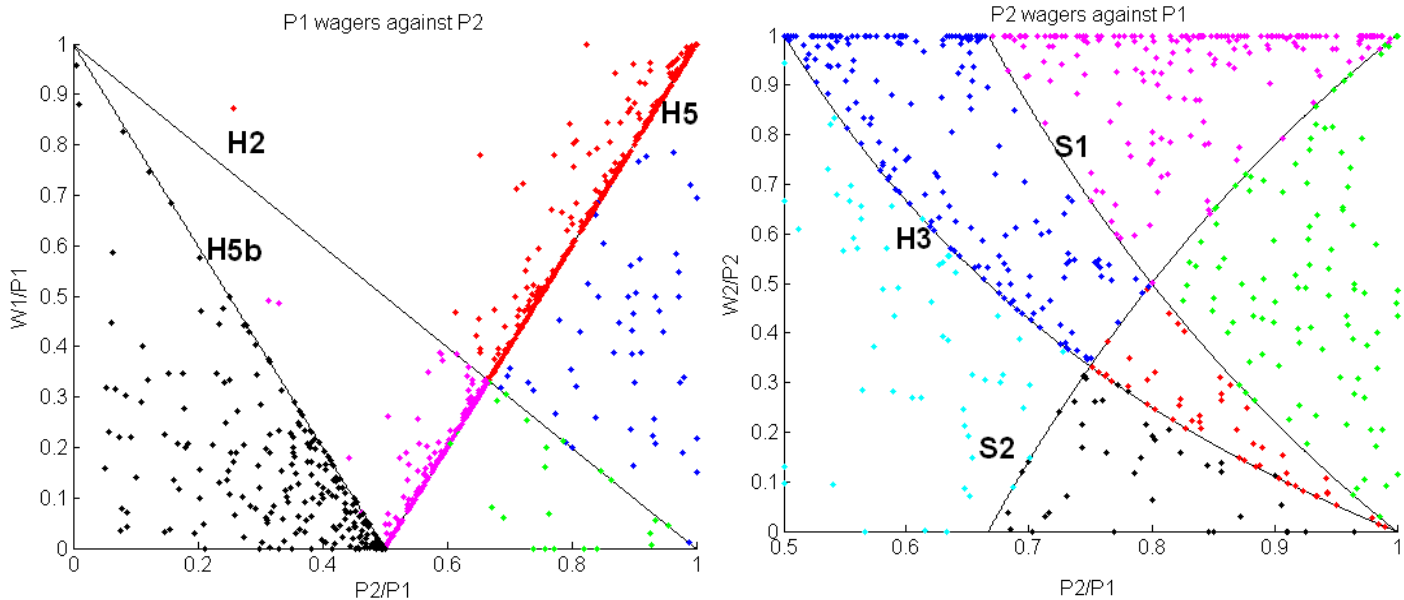
Figures 2 and 3: Player one wagering against Player two, and Player two wagering against Player one

There is a lot of information to unpack in the above figures. Most striking is how often (85%) player one employs the wager to cover player two doubling up (H5) or wagers right above it. Player two, on the other hand, most often (35%) wagers everything (or nearly so), or wagers just enough to get ahead of P_1 (H3, 14%). Although these can be good wagers, depending on the ratio of P_2 to P_1, they do not always make strategic sense. It appears as though adherence to the heuristics overwhelms rational optimization. This will be explored in the following sections. For Figure 3 it should be

noted that player two frequently wagers on heuristics related to player three, and these wagers will appear to be random when viewed just in the context of player two vs player one. The full paper explores this in more detail.

Player One overusing H5

While $P_2/P_1<0.5$ (case 1), player one can guarantee victory (R4) for himself as long as he wagers less than P_1-2*P_2 (H5b), although depending on his valuation of V, it can be rational to wager more. Similarly when $1/2<P_2/P_1<2/3$ (case 2) player one can guarantee R4 by wagering greater than $2*P_2-P_1$ (H5) and less than P_1-P_2 (H2). Therefore, in both cases player one has a minimum and maximum wager between which he has constant probability of victory. Where exactly the player decides to wager within these bounds should depend on his perceived odds of answering correctly and his risk aversion (it is impossible to separate these two factors, an issue explored in the paper), however the player should behave similarly in the two cases. In the first case, player one is guaranteed R4, the expected payout is $V+W_1*(Pr(1X)-Pr(0X))$, so a risk-neutral player should bet the maximum safe wager if he thinks he's more likely to get it right or not, wager nothing if he thinks otherwise and wager in between if he is unsure. Since the probability of player one answering correctly in this situation is 52%, the fact that the wagers represented by black dots in figure 2 are spread fairly evenly between the minimum and maximum boundaries implies that the player has a fairly accurate perception of the odds.

In the second case, $1/2<P_2/P_1<2/3$, the expected payout is $(P_1+V)*(Pr(1X)+Pr(00))+W_1*(Pr(1X)-Pr(00))+RU*Pr(01)$. A risk neutral player wanting to maximize the expected payout should maximize W_1 as long as he thinks the probability of answering correctly is greater than the probability of both himself and player two answering incorrectly. In this case the probability of getting it right is 48%, and both getting it wrong occurs only 29%. A player with reasonably accurate perceived odds should, therefore, bet at or near his maximum of H2. And yet figure 2 shows just the opposite behavior, most of the wagers are at the minimum along H5. Since there is no reason for players to so drastically shift their perceived odds of victory or risk behavior between cases one and two, it seems that the most likely explanation is that the H5 heuristic is very, very strong; players use it without realizing that they could win more by wagering up to H2 without decreasing their probability of winning. This is the most clean cut example in the data of a heuristic overwhelming predicted rational behavior.

Player two overusing H3

Similar to Player one in certain cases wagering H5 when it didn't make sense, Player two exhibits similar, although less drastic, behavior with H3. With $1/2<P_2/P_1<2/3$, any wager along H3 will put the outcome in R3 (assuming player one wagers above his minimum wager, which he always did in the games studied). But in R3 the only way for player two to win is to answer correctly while player one is incorrect. If player two is basing his wager solely around player one with no regard for player three, he might as well wager everything. This is not to say that player two should always wager everything in these cases, since it may be rational to wager to stay ahead of player three, depending on their respective scores. A similar case can be made for $2/3<P_2/P_1<3/4$, because in the case that H3 is above S2, meaning a wager along H3 will still put the game in R3 as long as player one is wagering at least H5. The point is that H3 is not strategically relevant in these cases, and yet as seen in figure 3, players use it as a guide just as frequently as when it is relevant in $P_2/P_1>3/4$.

Over-Confidence

Along with an over-reliance on the heuristics, the other factor leading to poor wagers appears to be overconfidence. For instance, player two wagers above S2 57% of the time, a wager that results in R3 if player one wagers at H5, which he does 85% of the time. This means player two is very frequently putting himself in a situation where he can only win if he gets it right and player one gets it wrong, which occurs fairly infrequently. The status quo shown in the data should not be an equilibrium; player two should be betting less much more often to take advantage of player one's over reliance on H5. And yet the data shows no trend of player two changing his behavior. One possibility is that players are over-confident. Even if they know the historic probabilities, they think that they are more likely to get the answer correct. If player two truly believes that he is very likely to get it right, then wagering nearly everything would be rational to him. This could explain why player two does this 35% of the time. The full paper provides a model for exploring how overconfident the players need to be to produce the wagers found in the data.

Conclusion

The purpose of this extended abstract was to introduce the basic framework used to analyze the actual wagering behavior of Jeopardy contestants, along with providing a few basic examples of how well known concepts from behavioral economics, heuristics and overconfidence, are causing seemingly irrational wagers. The full paper looks at the data in more detail and provides even more evidence to support these claims. It also includes the third player in the analysis, a factor largely ignored in this abstract. The goal of the full paper is to demonstrate that even something as trivial as a game show can provide real insights into human behavior.